



Vizalytics Knowledge Graph

A Vizalytics Technology White Paper

Abstract

The Vizalytics Knowledge Graph is a Software-as-a-Service mechanism that ingests event and logistical data from disparate sources, analyzes them according to time and location, generates contextual insights from relationships among the data sources, and notifies users in real time of impacts on their interests. The Knowledge Graph represents for its users what matters based on their intents and goals, why and how it matters, and when it matters – then suggests how to react. This is done automatically for different kinds of users – quickly, efficiently, and at scale. The Knowledge Graph is the foundation for a series of offerings from Vizalytics that generate location-specific business intelligence to help enterprises, governments, and private citizens derive value from private and open public data sets. The Vizalytics API can be tailored to add time-and-location-driven contextual analysis to corporate and institutional applications.

Creating Solutions

Vizalytics was born in the aftermath of Hurricane Sandy in 2012, along the U.S. eastern seaboard. The storm affected thousands of small businesses and shopkeepers and underscored two huge unmet challenges:

- The need for fast access to government data for shopkeepers struggling to survive and recover.
- The need for governments to learn how to analyze, coordinate, and distribute information from disparate data sources to their many constituencies.

But the need for a new data analytics tool went well beyond shopkeepers coping with a natural disaster. Day to day, local businesses need information to deal with fines and regulations, anticipate impacts on foot traffic, and become aware of potential business disruptions. Governments and their agencies cope with stovepiped organizations, decreasing budgets and staff, and limited transparency and data normalization. Enterprises swim in too much content that is unorganized and hard to find.

Thus a software solution to a data distribution problem needed to push beyond simple access:

- Available data comes from a variety of government and private sources.
- The databases are often unwieldy if they are accessible at all.
- Often to work with this data and draw meaningful conclusions requires expertise not available to potential users – be they storeowners, citizens, or government employees.
- The challenge is to go beyond mere filtering of individual data sets to find relationships among the sources and tease meaningful insights from those relationships that are relevant to particular user goals.
- The software must normalize reports from the source data sets in terms of nomenclature and display.
- To be truly useful actionable insights must be available in real time.





A given organization might have dozens of distinct data streams to deal with – inspections, traffic, transit, construction, health department, permitting, crime, census, remote sensors. Each stream monitors its own discrete set of events, often hundreds of them, in granular detail. At a given time, events can affect activities in a given area interdependently and differently depending on context – when they occur, where they occur, and what the activity is. Certain events are more, less, or differently impactful in the presence or absence of other events.

Having advance or real-time notice of particular events, and understanding their interaction probabilities or impacts, allows for proactive response to mitigate or leverage them:

- Store owners made aware of unsanitary conditions at or near their place of business can take action to avoid fines or lost business.
- Prior knowledge of sidewalk or street closures that might reduce foot traffic to a business can enable the owner to make decisions about staff deployment, signage, or other actions to mitigate the impact.
- Timely awareness of criminal activity enables a business owner to bolster security or warn customers, or a resident to counsel children.
- A notice about street closures or transit delays might cause a tourist to change plans to avoid inconvenience.

How to create a Software-as-a-Service solution for this multi-dimensional problem? And how widely might this contextual analysis apply in business and government?

Semantic Analysis

The Vizalytics Knowledge Graph is a living, continuously updating and evolving spatio-temporal knowledge base that maps relationships among data sets in real time. It was created to break down walls between disparate data streams, normalize nomenclature and display, and map the events – upwards of 500 distinct event types for a large urban application – and relationships among them against time and location. Then, using context-driven analytical “lenses,” the Knowledge Graph returns facts appropriate to the requirements of the requester. The process is a combination of location, time, event, insight, and device.

Level 1: Facts

Knowledge Graph analytics begin with facts drawn from two types of data relevant to a particular geography – say a city, a neighborhood, a building, a manufacturing plant, or a vehicle:

- Semi-static Data – the sort of data that would populate a GIS layer such as street addresses, building characteristics, businesses in a building, number of floors, residential/commercial mix, utility installations, mass transit access, distance to parks or schools, topography, real estate information, census data, and the like.
- Dynamic Data – event data such as street closings, 311 requests, inspections, crime reports, transit or traffic changes or delays, scheduled activities, and the like.





Level 2: Analytics

The cloud-based Knowledge Graph creates a spatio-temporal model of a given geography at hyperlocal granularity. Semantic analysis extracts meaning from disparate data facts as they apply to the attributes of each user. Through dynamic impact analysis and relevance-ranking algorithms, the Knowledge Graph uses machine learning, graph analysis, fuzzy logic, and ontology-driven inference to assess significance of events, whether they are typical or unusual, where related events are clustered, their historical context, and what impact they may have on a user.

Level 3: Impacts

Based on its assessment of the data-supplied facts, the Knowledge Graph returns insights, predictions, risks, and scores (such as business viability or neighborhood ambiance score – “your business has an 87% chance of having a health department inspection”). Based on a user’s location, the context of a notification may vary: A rodent was reported next door to your business, around the corner, or in your neighborhood, depending on the governing geographic area – a specified polygon on a map.

Level 4: Content

Using natural language generation techniques, the Knowledge Graph delivers consumable insights and impact statements in response to user goals, preferences, and characteristics. Along with impact-related insights, the Knowledge Graph – based on triggers aligned with user needs and characteristics – provides context-sensitive tips and suggestions that are not precisely driven by fine-grained events:

- “This regulation changed and you may be affected by it.”
- “Here’s a new business course that might be relevant to your enterprise.”
- “This news article reports on construction in your neighborhood.”

The data and content may be delivered as rolled-up statements or as data points suitable for an analytics dashboard on a desktop or mobile device.

Offerings from the Knowledge Graph

Building on the Knowledge Graph’s automated, scalable insight pipeline, Vizalytics customizes open data solutions via five basic offers:

- *Vizalytics Business Builder*, which uses data-driven, real-time competitive rankings, feasibility scores, and customer/neighborhood profiling to speed site selection, location performance, use case optimization, development, and future planning decisions.
- *Vizalytics Local Insight*, an enterprise-level dashboard for governments and enterprises. Local Insight aggregates information from multiple sources of governmental and public data, generating custom insights for each user, producing a multi-layered view of conditions, events, and changes within a specified geography. A total Local Insight solution can include Mind My Business™ and “neighborhoods” elements as required.





- *The Vizalytics Application Programming Interface (API)* gives customers access to the Vizalytics Knowledge Graph to integrate contextual intelligence analytics into their own applications.
- *The Vizalytics Mind My Business™ mobile app for shopkeepers.* We alert users about construction, traffic, regulatory issues, health and safety concerns, fines, events, and other information. The app helps shopkeepers make more informed decisions, creating opportunities and saving them time and money while improving regulatory compliance.
- *A Vizalytics “neighborhoods” website,* a resource for residents and a catalyst for community organizing and change. In New York City, the neighborhoods.nyc site provides hyper-local traffic, transit, quality of life, health, inspection, event, and other information, plus a connection to city services for more than 400 neighborhoods.

Knowledge Graph Use Cases

- A bank might apply a Mind My Business™ model to its mobile banking application, or a business prospecting service for loan customers.
- A news organization might add a “neighborhoods” component to its public web site, while using a Local Insight platform to support news gathering.
- A real estate company might apply Knowledge Graph analytics to building management or development.

The Vizalytics Knowledge Graph defines and represents an intent proposition from a user in an abstract way, then generates relevant contextual insight for different classes of users with different needs – **and does so at scale**. These insights are based on impactful events that have occurred, are occurring, or are expected to occur, in proximity to the user’s location of interest. The system computes the potential impact of these events on the user’s goals at the time of inquiry, **representing a level of reality that human senses cannot perceive**.

In this way a collection of disconnected data sources becomes, through the Knowledge Graph, a unified and continuously-evolving portal to fast-changing environments.

Vizalytics Technology
40 Wall 28th Fl NY NY 10005
info@vizalytics.com
Vizalytics.com

Aileen Gemma Smith, CEO
aileen@vizalytics.com
+1 (718) 873-6548

